Docket No. AUS920000843US1

APPARATUS, METHOD AND COMPUTER PROGRAM PRODUCT FOR STOPPING PROCESSORS WITHOUT USING NON-MASKABLE INTERRUPTS

BACKGROUND OF THE INVENTION

5

1.   Technical Field:

The present invention is directed to an apparatus, method, and computer program product for stopping processors without using non-maskable interrupts. More
10 specifically, the present invention is direct to an apparatus, method and computer program product for stopping processors in a multiprocessor system that have interrupts disabled without using non-maskable interrupts.

15

2.   Description of Related Art:

In multiprocessor systems, there are times when a processor in the multiprocessor system experiences an event, or the debugger is entered, yet the software is
20 unable to stop all of the other processors. An "event" is any occurrence that causes one of the processors to enter a debugger, e.g. an error, an exception, an interrupt, or the like. This is usually due to the other processors looping in a portion of code that is in a
25 disabled state, i.e. code that has disabled interrupts on the processor.

This situation has typically been handled by the use of interrupts. An interrupt is a signal informing a program that an event has occurred. When a program
30 receives an interrupt signal, it takes a specified action (which can be to ignore the signal). Interrupt signals

can cause a program to suspend itself temporarily to

service the interrupt.

Interrupt signals can come from a variety of sources. For example, every keystroke on a keyboard generates an interrupt signal. Interrupts can also be
5 generated by other devices, such as a printer, to indicate that some event has occurred. These are called hardware interrupts.

Interrupt signals initiated by programs are called software interrupts. A software interrupt is also called
10 a trap or an exception. Each type of software interrupt is associated with an interrupt handler, i.e. a routine that takes control when the interrupt occurs. For example, when a key is pressed on a keyboard, this action triggers a specific interrupt handler. The complete list
15 of interrupts and associated interrupt handlers is stored in a table called the interrupt vector table.

The particular interrupt used to address the problem described above with regard to multiprocessor systems is called a non-maskable interrupt. A non-maskable
20 interrupt (NMI) is a high-priority interrupt that cannot be disabled by another interrupt. NMIs are used to report malfunctions such as parity, bus and math coprocessor errors. The NMI is non-maskable in that this interrupt can not be disabled by software and cannot be
25 ignored by the system.

Not all multiprocessor systems support NMI and thus, not all multiprocessor systems are capable of using NMI to stop processors that are looping in disabled code when an event on one of the processors occurs. Moreover, even
30 on multiprocessor systems that support an NMI mechanism, asserting the NMI from software in a recoverable way in order to stop other processors in the system is not a

Docket No. AUS920000843US1

simple task. Thus, it would be beneficial to have an apparatus, method and computer program product for stopping processors in a multiprocessor system without using NMI.

SUMMARY OF THE INVENTION

The present invention provides an apparatus, method and computer program product for stopping processors in a multiprocessor system without using non-maskable interrupts. With the apparatus, method and computer program product of the present invention, at system initialization time, a copy of the operating system (OS) kernel is copied to a new physical location in memory. When a processor enters the debugger due to the occurrence of an event, such as encountering a breakpoint, a trigger, a watchpoint, the occurrence of an error, or the like, the debugger switches its virtual-to-physical address mapping to point to the new copy of the OS kernel. The original copy of the OS kernel is then modified by inserting breakpoints, e.g., interrupts, in a repeating pattern in the text of the original copy of the OS kernel, with the exception of the breakpoint handler text in the original copy of the OS kernel.

The debugger then performs architecture dependent actions to flush the executing processors' data caches so that the modifications are present in memory and therefore, visible to the instruction stream of the other processors. The debugger then performs architecture dependent actions to broadcast instruction cache invalidate operations to thereby force the processors to refetch instructions from the OS kernel.

When the remaining processors fetch the OS kernel instructions, the instructions are fetched from the modified OS kernel. Thus, the processors encounter the

Docket No. AUS920000843US1

inserted breakpoints and enter a breakpoint handler. The breakpoint handler then, by virtue of the switched virtual-to-physical address mapping, redirects the processor to the new copy of the OS kernel and handles

5   the breakpoint in a normal fashion, e.g. causes the processor to enter the debugger. The debugger, at this point, now has control over all of the processors in the multiprocessor system. Thus, the debugger is now able to diagnose what the other processors in the multiprocessor

10 system were doing at the time the event occurred. Once the event is handled, i.e. the debugger is exited, the original OS kernel may be recovered from the new copy, the virtual-to-physical mapping may be reset to its original state, and the system may be recovered.

15

Docket No. AUS920000843US1

BRIEF DESCRIPTION OF THE DRAWINGS

The novel features believed characteristic of the invention are set forth in the appended claims. The
5 invention itself, however, as well as a preferred mode of use, further objectives and advantages thereof, will best be understood by reference to the following detailed description of an illustrative embodiment when read in conjunction with the accompanying drawings, wherein:

10 **Figure 1** is an exemplary block diagram of a data processing system in accordance with the present invention;

Figure 2 is a diagram illustrating the operation of the data processing system when the present invention is
15 employed; and

Figure 3 is a flowchart outlining an exemplary operation of the present invention.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

With reference now to the figures, **Figure 1** is a
block diagram of a multiprocessor data processing system
5   in which the present invention may be employed. As shown
in **Figure 1**, the data processing system **100** may be a
symmetric multiprocessor (SMP) system, a Non-Uniform
Memory Access (NUMA) system, or other multiprocessor
system that includes a plurality of processors **102-106**
10   connected to system bus **107**. Also connected to system
bus **107** is memory controller/cache **108**, which provides an
interface to local memory **109**. I/O bus bridge **110** is
connected to system bus **107** and provides an interface to
I/O bus **112**. Memory controller/cache **108** and I/O bus
15   bridge **110** may be integrated as depicted.

Peripheral component interconnect (PCI) bus bridge
114 connected to I/O bus **112** provides an interface to PCI
local bus **116**. A number of modems may be connected to
PCI bus **116**. Typical PCI bus implementations will
20   support four PCI expansion slots or add-in connectors.
Communications links to other devices, networks, and the
like, may be provided through modem **118** and network
adapter **120** connected to PCI local bus **116** through add-in
boards.

25   Additional PCI bus bridges **122** and **124** provide
interfaces for additional PCI buses **126** and **128**, from
which additional modems or network adapters may be
supported. In this manner, data processing system **100**
allows connections to multiple network computers. A
30   memory-mapped graphics adapter **130** and hard disk **132** may

also be connected to I/O bus **112** as depicted, either

directly or indirectly.

Those of ordinary skill in the art will appreciate that the hardware depicted in **Figure 1** may vary. For example, other peripheral devices, such as optical disk drives and the like, also may be used in addition to or in place of the hardware depicted. Moreover, many of the elements shown in **Figure 1** may not be present in the data processing system in which the present invention is employed. The depicted example is not meant to imply architectural limitations with respect to the present invention.

The data processing system depicted in **Figure 1** may be, for example, an IBM RISC/System 6000 system, a product of International Business Machines Corporation in Armonk, New York, running the Advanced Interactive Executive (AIX) operating system.

The present invention may be implemented in hardware, software, or a combination of hardware and software within the multiprocessor system. For example, the present invention may be implemented as a hardware based debugger element in the multiprocessor system, as instructions executed by one or more of the processors in the multiprocessor system, or the like. Some of the beneficial aspects of the present invention are that the present invention does not require external hardware or internal hardware support, can work across the spectrum of known computer systems (whether they support NMI or not), and allows a self-hosted debugger to perform the functions of the present invention.

Figure 2 is a diagram illustrating the operation of the data processing system of **Figure 1** when the present invention is employed. As shown in **Figure 2**, the operation starts with initialization of the system at

Docket No. AUS920000843US1

which time the OS kernel is copied to a new physical location (**210**). The copying of the OS kernel and debugger in a preferred embodiment, involves copying the OS kernel and debugger text using text replication.

5   However, other mechanisms for copying the OS kernel to a new physical location may be utilized without departing from the spirit and scope of the present invention. Moreover, the copying of the OS kernel may be performed at other times than at system initialization, such as,
10   for example, when an event occurs and the debugger is entered.

At some time later, an event occurs in one or more of the processors of the multiprocessor data processing system (**220**). An "event" is any occurrence that causes
15   one of the processors in the multiprocessor system to enter a debugger, e.g. an error, an exception, an interrupt, or the like. The debugger for the processor is entered (**230**). The debugger switches the virtual-to-physical address mapping of the processor so
20   that the mapping now points to the new copy of the OS kernel (**240**). The virtual-to-physical switching is a processor architecture dependent action which translates to reinitializing page table entries or translation lookaside buffers (TLBs). It involves pointing the
25   virtual addresses to different physical addresses.

The debugger then inserts breakpoints, e.g., interrupts, in the original copy of the OS kernel (**250**). The insertion of the breakpoints, in a preferred embodiment, involves inserting a repeating pattern of
30   architecture dependent breakpoint instruction (also referred to as "trap" instruction or "interrupt" instruction) opcodes into all portions of the OS kernel with the exception of the breakpoint handler. This is

because the breakpoint handler must continue to be able to handle breakpoints as they occur.

　　　For example, a breakpoint may instruct a processor to throw an exception causing the processor to initiate a

5　debugger. The repeating pattern may be such that, for example, every other instruction is a breakpoint instruction, thereby assuring that no matter where in the OS kernel a processor retrieves instructions, at least one breakpoint instruction will be encountered.

10　　　The debugger then performs architecture specific functions to flush the caches of the other processors in the multiprocessor system to thereby force the processors to re-fetch the instructions which they were executing (**260**) from the OS kernel. Such architecture specific

15　functions to flush the caches may include, for example, issuing a data cache block flush instruction (dcbf) in a PowerPC architecture or a flush cache (fc) instruction in an IA64 architecture. Moreover, an instruction to invalidate the instruction cache may also be used, such

20　as instruction cache block invalidate (icbi) in the PowerPC architecture or the flush cache (fc) instruction in the IA64 architecture.

　　　Once the processor caches are flushed, the processors then refetch their instructions from the OS

25　kernel (**270**) and encounter the inserted breakpoints (**280**). The breakpoints cause the processors to enter the breakpoint handler of the modified OS kernel (**290**). Once the processors enter the breakpoint handler of the modified OS kernel, the virtual-to-physical mapping of

30　the processor is switched to point to the new copy of the OS kernel (**292**). Such switching is facilitated by code in the breakpoint handler that identifies when one or more other processors in the multiprocessor system have

Docket No. AUS920000843US1

switched to a different copy of the OS kernel. Such
identification may make use of a global flag, for
example. The new copy of the OS kernel then causes the
breakpoint handler to instruct the processor to enter the
5   debugger (**295**).

As an example of the functioning of the present
invention, the physical address 0x2000 may contain an
original copy of kernel while the physical address 0x8000
may contain a new copy of kernel. For the normal running
10 kernel, virtual address 0xC0000000 is mapped to physical
address 0x2000, and this mapping is true for all
processors in the system. Thus, every processor uses
virtual address 0xC0000000 to access kernel.

When the first processor enters the debugger, this
15 processor switches its translation hardware so that its
virtual address 0xC0000000 now maps to physical address
0x8000, i.e. The new copy of the OS kernel. All
processors are still running the kernel at virtual
address 0xC0000000, but now one of them is actually using
20 a different physical copy.

After the debugger inserts breakpoints in the
original physical copy of the OS kernel (at physical
address 0x2000), the other processors encounter the
breakpoints and enter the breakpoint handler (a portion
25 of the 0x2000 text that was not modified to include
breakpoints). The code in the breakpoint handler
recognizes (via a global flag) that another processor has
switched to a new copy of OS kernel, and the other
processors do the same, i.e. they change their
30 translation hardware to map 0xC0000000 to 0x8000. The
processors are then all running the new copy of the OS
kernel.

In this way, the debugger now gains control over

each of the processors in the multiprocessor system, i.e. the looping of the processors is stopped by entering the debugger. Some of these processors may have been in a fully disabled state and the debugger would not have been

5   able to gain control over these processors in the prior art systems without the benefit of the present invention.

Figure 3 is a flowchart outlining an exemplary operation of the present invention. While the operational steps shown in **Figure 3** are shown in an

10 exemplary order, the order of certain steps in the operation in **Figure 3** may be rearranged without departing from the spirit and scope of the present invention. For example, steps **310** and **320**, as will be described hereafter, may be swapped such that copying of the OS

15 kernel may be performed after entry into the debugger.

As shown in **Figure 3**, the operation starts with copying the OS kernel to a new physical location (step **310**). An event occurs and the debugger is entered (step **320**). The virtual-to-physical address mapping of the

20 processors is switched to point to the new copy of the OS kernel (step **330**). Breakpoints are inserted into the original copy of the OS kernel in all portions with the exception of the breakpoint handler (step **340**). Actions to flush the executing processors' data caches are

25 performed (step **350**) and the operation ends. By virtue of this operation, the processors then refetch the instructions from the modified OS kernel, encounter a breakpoint, enter the breakpoint handler of the OS kernel and are redirected to the debugger.

30   Thus, the present invention provides a mechanism by which all of the processors in a multiprocessor system may be stopped when there is an occurrence of an event

Docket No. AUS920000843US1

that causes entry into a debugger. With the present
invention, even processors that are in a fully disabled
state, i.e. are looping on disabled instructions, may be
stopped using the present invention. Furthermore, the

5   present invention may be used regardless of whether the
multiprocessor system supports NMIs. Moreover, the
present invention provides a mechanism by which all of
the processors in a multiprocessor system may be stopped
and recovery of the system may be performed with ease.

10   It is important to note that while the present
invention has been described in the context of a fully
functioning data processing system, those of ordinary
skill in the art will appreciate that the processes of
the present invention are capable of being distributed in

15   the form of a computer readable medium of instructions
and a variety of forms and that the present invention
applies equally regardless of the particular type of
signal bearing media actually used to carry out the
distribution. Examples of computer readable media

20   include recordable-type media such a floppy disc, a hard
disk drive, a RAM, and CD-ROMs and transmission-type
media such as digital and analog communications links.

        The description of the present invention has been
presented for purposes of illustration and description,

25   but is not intended to be exhaustive or limited to the
invention in the form disclosed. Many modifications and
variations will be apparent to those of ordinary skill in
the art. The embodiment was chosen and described in
order to best explain the principles of the invention,

30   the practical application, and to enable others of
ordinary skill in the art to understand the invention for
various embodiments with various modifications as are
suited to the particular use contemplated.